

CS 547 Lecture 10: The Poisson Process

Daniel Myers

Up to this point, we've discussed arrivals to queueing systems in a general way, and we've used the general arrival rate λ as part of our asymptotic bounds calculations. Now, we'll actually describe the statistical behavior of one particular arrival process: the Poisson process. Poisson arrivals are by far the most popular arrival model used in the analysis of queueing systems.

Counting Events

The Poisson process describes the statistical properties of a sequence of *events*. In our case, these events will usually be arrivals to a queueing system, but other types of events could be used in other applications.

Let $N(t)$ represent the number of events that occur in the interval $[0, t]$.

Little-o Notation

We say that a function $f(h)$ is $o(h)$ if $f(h)$ goes to zero faster than h . That is,

$$\lim_{h \rightarrow 0} \frac{f(h)}{h} = 0$$

Definition of the Poisson Process

The sequence of random variables $\{N(t), t \geq 0\}$ is said to be a *Poisson process with rate $\lambda > 0$* if the following five conditions hold.

1. $N(0) = 0$
2. The numbers of events that occur in non-overlapping time periods are independent
3. The distribution of the number of events that occur in a given period depends only on the length of the period and not its location
4. $P\{N(h) = 1\} = \lambda h + o(h)$
5. $P\{N(h) \geq 2\} = o(h)$

The first property establishes the initial condition: at time 0, no events have occurred.

The second property says that we can treat disjoint regions as statistically independent. That is, the number of events that occur in one period of time has no influence on the number of events that occur in a different period of time. This should seem somewhat reminiscent of the memoryless property. We'll use this property several times in our derivations.

The third property states that the process is stable over time. That is, the statistical properties of the process do not change as time advances, so we can regard all periods of the same length as statistically identical, regardless of where in time the period actually takes place.

The fourth and fifth properties tell us about the behavior of the process in very small windows of time. The fourth property says that the probability of getting one event in a small window of time is proportional to the length of the window, plus some error term that goes to zero as the window size shrinks to zero. The fifth property says that the probability of getting two or more events in a window of time goes to zero as the window becomes very small.

Using the fourth and fifth properties, we can derive a simple proposition.

$$\begin{aligned} P\{N(h) = 0\} &= 1 - P\{N(h) \geq 1\} \\ &= 1 - \lambda h - o(h) \end{aligned}$$

Key Properties of the Poisson Process

Using the definition of the Poisson process, we can now derive three key results that summarize its behavior.

- The number of events in a period is Poisson
- Events in the process occur at constant rate λ
- The time between events is exponentially distributed with parameter λ

The Distribution of the Number of Events

In this section, we'll show that the number of events generated by the Poisson process over a period of length t follows a Poisson distribution.

First, we'll derive an expression for $P\{N(t) = 0\}$, then state the general distribution.

Consider $P\{N(t+h) = 0\}$. If we get no events in $[0, t+h]$, then we received no events in $[0, t]$ and no events in $[t, t+h]$. Therefore,

$$P\{N(t+h) = 0\} = P\{N(t) = 0 \text{ and } N(t+h) - N(t) = 0\}$$

The two periods are non-overlapping, so we can treat them independently (using the second part of the definition), and write the joint probability as the product of the individual probabilities.

$$P\{N(t+h) = 0\} = P\{N(t) = 0\} P\{N(t+h) - N(t) = 0\}$$

The second term is simply a period of length h , which is statistically identical all other periods of length h by the third part of the definition of the Poisson process.

$$\begin{aligned} P\{N(t+h) = 0\} &= P\{N(t) = 0\} P\{N(h) = 0\} \\ &= P\{N(t) = 0\} (1 - \lambda h - o(h)) \end{aligned}$$

To simplify notation, let $P_0(t) = \{N(t) = 0\}$. We now have

$$P_0(t+h) = P_0(t)(1 - \lambda h - o(h))$$

Rearranging and dividing both sides by h ,

$$\frac{P_0(t+h) - P_0(t)}{h} = -\lambda P_0(t) - \frac{o(h)}{h} P_0(t)$$

Now, take the limit of both sides as $h \rightarrow 0$.

$$\lim_{h \rightarrow 0} \frac{P_0(t+h) - P_0(t)}{h} = \lim_{h \rightarrow 0} -\lambda P_0(t) - \frac{o(h)}{h} P_0(t)$$

The $o(h)$ term goes to zero. The left hand side is simply the definition of the derivative of $P_0(t)$, so we obtain a differential equation.

$$\frac{\partial}{\partial t} P_0(t) = -\lambda P_0(t)$$

Solving the equation yields

$$P_0(t) = P\{N(t) = 0\} = e^{-\lambda t}$$

This result can be extended to show that the probability of getting k events during a period of length t follows a Poisson distribution with parameter λ .

$$P\{N(t) = k\} = \frac{e^{-\lambda t} (\lambda t)^k}{k!}$$

Setting $k = 0$ gives $e^{-\lambda t}$, as expected.

The expected number of events in a period is simply

$$E[N(t)] = \lambda t$$

Events Occur At Constant Rate

By definition, the rate of arrivals is the expected number of arrivals in a period, divided by the length of the period.

$$\text{rate of arrivals} = \frac{\text{expected number of arrivals}}{\text{length of period}}$$

For the Poisson process, the rate of arrivals over a period of length t is

$$\text{rate of arrivals over } t = \frac{\lambda t}{t} = \lambda$$

The expected rate of arrivals over any period is simply λ , regardless of the length of the period.

By extension, the expected time between arrivals is $\frac{1}{\lambda}$.

Note: a constant arrival rate is not the same as a constant interarrival time!

The Time Between Events is Exponentially Distributed

Let T_1 be the time until the first event, and T_2 be the time between the first and second events. We would like to derive the distribution of T_2 , and by extension, the distribution of the time between any two events.

Suppose the first event occurs at time t . We'll reason about the CCDF of T_2 conditioned on $T_1 = t$.

$$P\{T_2 > s \mid T_1 = t\}$$

T_2 is the time between the first and second events, so if $T_2 > s$, no events can occur in the time between t and $t + s$.

$$P\{T_2 > s \mid T_1 = t\} = P\{\text{no events in } [t, t + s] \mid T_1 = t\}$$

The time periods $[0, t]$ and $[t, t + s]$ are non-overlapping, so we can treat them independently. Therefore, we can remove the conditioning on T_1 .

$$P\{T_2 > s \mid T_1 = t\} = P\{\text{no events in } [t, t + s]\}$$

Further, the period from t to $t + s$ is simply a period of length s , so we can treat it the same as any other period of length s .

$$\begin{aligned} P\{T_2 > s \mid T_1 = t\} &= P\{\text{no events in a period of length } s\} \\ &= P\{N(s) = 0\} \\ &= e^{-\lambda s} \end{aligned}$$

The final step is simply our previous result for the probability of getting 0 events during a period.

Therefore, T_2 has the CCDF of an exponential distribution with parameter λ .